

# Evaluation of Labeling Strategies for Rotating Maps

Andreas Gemsa, Martin Nöllenburg, and Ignaz Rutter

Institute of Theoretical Informatics, Karlsruhe Institute of Technology,  
Am Fasanengarten 5, 76131 Karlsruhe, Germany

**Abstract.** We consider the following problem of labeling points in a dynamic map that allows rotation. We are given a set of points in the plane labeled by a set of mutually disjoint labels, where each label is an axis-aligned rectangle attached with one corner to its respective point. We require that each label remains horizontally aligned during the map rotation and our goal is to find a set of mutually non-overlapping *active* labels for every rotation angle  $\alpha \in [0, 2\pi)$  so that the number of active labels over a full map rotation of  $2\pi$  is maximized.

We discuss and experimentally evaluate several labeling models that define additional consistency constraints on label activities in order to reduce flickering effects during monotone map rotation. We introduce three heuristic algorithms and compare them experimentally to an existing approximation algorithm and exact solutions obtained from an integer linear program. Our results show that on the one hand low flickering can be achieved at the expense of only a small reduction in the objective value, and that on the other hand the proposed heuristics achieve a high labeling quality significantly faster than the other methods.

## 1 Introduction

Dynamic digital maps, in which users can navigate by continuously zooming, panning, or rotating their personal map view, opened up a new era in cartography and geographic information science (GIS) from professional applications to personal mapping services on mobile devices. The continuously animated map view adds a temporal dimension to the map layout and thus many traditional algorithms for static maps do not extend easily to dynamic maps. Despite the popularity and widespread use of dynamic maps, relatively little attention has been paid to provably good or experimentally evaluated algorithms for dynamic maps.

In this paper we consider *dynamic map labeling* for points, i.e., the problem of deciding when and where to show labels for a set of point features on a map in such a way that visually distracting effects during map animation are kept to a minimum. In particular, we study rotating maps, where the mode of interaction is restricted to changing the map orientation, e.g., to be aligned with the travel direction in a car navigation system.

Been et al. [2, 3] defined a set of *consistency desiderata* for labeling zoomable dynamic maps, which include that (i) labels do not *pop* or *flicker* during monotone zooming, (ii) labels do not *jump* during the animation, and (iii) the labeling only depends on the current view and not its history. In our previous paper [8], we adapted the consistency model of Been et al. to rotating maps, showed NP-hardness and other properties of consistent labelings in this model, and provided efficient approximation algorithms.

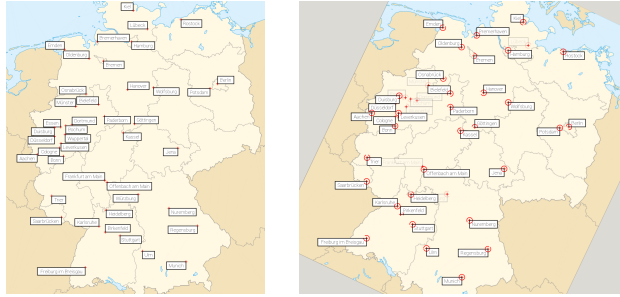


Fig. 1: Instance with 43 labeled cities in Germany. Input labeling (left), rotated by  $\sim 25^\circ$  (right). Background picture is in public domain. Retrieved from Wikipedia [Link]

Similar to the (NP-hard) label number maximization problem in static map labeling [6], the goal in dynamic map labeling is to maximize the number of visible or *active* labels integrated over one full rotation of  $2\pi$ . The value of this integral is denoted as the *total activity* and defines our objective function. Figure 1 shows an example seen from two different angles. Without any consistency restrictions, we can select the active labels for every rotation angle  $\alpha \in [0, 2\pi)$  independently of any other rotation angles. Clearly, this may produce an arbitrarily high number of flickering effects that occur whenever a label changes from active to inactive or vice versa. Depending on the actual consistency model, the number of flickering events per label is usually restricted to a very small number. Our goal in this paper is to evaluate several possible labeling strategies, where a labeling strategy combines both a consistency model and a labeling algorithm. First, we want to evaluate the loss in total activity caused by using a specific consistent labeling model rather than an unrestricted one. Second, we are interested in evaluating how close to the optimum total activity our proposed algorithms get for real-world instances in a given consistency model.

*Related Work.* Most previous work on dynamic map labeling covers maps that allow panning and zooming, e.g., [2, 3, 11, 12, 14]; there is also some work on labeling dynamic points in a static map [4, 5]. As mentioned above, the dynamic map labeling problem for rotating maps has first been considered in our previous paper [8]. We introduced a consistency model, and proved NP-completeness even for unit-square labels. For unit-height labels we described an efficient  $\frac{1}{4}$ -approximation algorithm as well as a PTAS. Yokosuka and Imai [15] considered the label size maximization problem for rotating maps, where the goal is to find the maximum font size for which all labels can be constantly active during rotation. Finally Gemsa et al. [7] studied a trajectory-based labeling model, in which a locally consistent labeling for a viewport moving along a given smooth trajectory needs to be computed. Their model combines panning and rotation of the map view.

*Our Contribution.* In this paper we take a practical point of view on the dynamic map labeling problem for rotating maps. In Section 2 we formally introduce the algorithmic problem and discuss our original rather strict consistency model [8], as well as two possible relaxations that are interesting in practice. Section 3 introduces three greedy heuristics (one of which is a  $\frac{1}{8}$ -approximation for unit square labels) and presents a for-

mulation as an integer linear program (ILP), which provides us with optimal solutions against which to compare the algorithms. Our main contribution is the experimental evaluation in Section 4. We extracted several real-world labeling instances from OpenStreetMap data and make them available as a benchmark set. Based on these data, we evaluate both the trade-off between the consistency and the total activity, and the performance of the proposed labeling algorithms. The experimental results indicate that a high degree of labeling consistency can be obtained at a very small loss in activity. Moreover, our greedy algorithms achieve a high labeling quality and outperform the running times of the other methods by several orders of magnitude. We conclude with a suggestion of the most promising labeling strategies for typical use cases.

Due to space constraints we omitted all proofs in this paper. They can be found, together with a more detailed experimental analysis, in the full version [9].

## 2 Preliminaries

In this section we describe a general labeling model for rotating maps with axis-aligned rectangular labels. This model extends our earlier model [8].

Let  $M$  be an (abstract) map, consisting of a set  $P = \{p_1, \dots, p_n\}$  of points in the plane together with a set  $L = \{\ell_1, \dots, \ell_n\}$  of pairwise disjoint, closed, and axis-aligned rectangular labels in the plane. Each point  $p_i$  must coincide with a corner of its corresponding label  $\ell_i$ ; we denote that corner (and the point  $p_i$ ) as the *anchor* of label  $\ell_i$ . Since each label has four possible positions with respect to  $p_i$  this widely used model is known in the literature as the 4-position model (4P) [6].

As  $M$  rotates, each label  $\ell_i$  in  $L$  must remain horizontally aligned and anchored at  $p_i$ . Thus, new label intersections form and existing ones disappear during the rotation of  $M$ . We take the following alternative perspective on the rotation of  $M$ . Rather than rotating the points, say clockwise, and keeping the labels horizontally aligned we may instead rotate each label counterclockwise around its anchor point and keep the set of points fixed. Both rotations are equivalent in the sense that they yield exactly the same intersections of labels and occlusions of points.

We consider all rotation angles modulo  $2\pi$ . For convenience we introduce the interval notation  $[a, b]$  for any two angles  $a, b \in [0, 2\pi]$ . If  $a \leq b$ , this corresponds to the standard meaning of an interval, otherwise, if  $a > b$ , we define  $[a, b] := [a, 2\pi] \cup [0, b]$ . For simplicity, we refer to any set of the form  $[a, b]$  as an interval. We define the length of an interval  $I = [a, b]$  as  $|I| = b - a$  if  $a \leq b$  and  $|I| = 2\pi - a + b$  if  $a > b$ .

A *rotation* of  $L$  is defined by a rotation angle  $\alpha \in [0, 2\pi)$ . We define  $L(\alpha)$  as the set of all labels, each rotated by an angle of  $\alpha$  around its anchor point. A *rotation labeling* of  $M$  is a function  $\phi: L \times [0, 2\pi) \rightarrow \{0, 1\}$  such that  $\phi(\ell, \alpha) = 1$  if label  $\ell$  is visible or *active* in the rotation of  $L$  by  $\alpha$ , and  $\phi(\ell, \alpha) = 0$  otherwise. We call a labeling  $\phi$  *valid* if, for any rotation  $\alpha$ , the set of labels  $L_\phi(\alpha) = \{\ell \in L(\alpha) \mid \phi(\ell, \alpha) = 1\}$  consists of pairwise disjoint labels. If two labels  $\ell$  and  $\ell'$  in  $L(\alpha)$  intersect, we say that they have a (soft) *conflict* at  $\alpha$ , i.e., in a valid labeling at most one of them can be active at  $\alpha$ . We define the set  $C(\ell, \ell') = \{\alpha \in [0, 2\pi) \mid \ell \text{ and } \ell' \text{ are in conflict at } \alpha\}$  as the *conflict set* of  $\ell$  and  $\ell'$ . Further, we call a contiguous range in  $C(\ell, \ell')$  a *conflict range*. The begin and end of a maximal conflict range are called *conflict events*.

For a label  $\ell$  we call each maximal interval  $I \subseteq [0, 2\pi)$  with  $\phi(\ell, \alpha) = 1$  for all  $\alpha \in I$  an *active range* of label  $\ell$  and define the set  $A_\phi(\ell)$  as the set of all active ranges of  $\ell$  in  $\phi$ . We call an active range where both boundaries are conflict events a *regular active range*. Our optimization goal is to find a valid labeling  $\phi$  that shows a maximum number of labels integrated over one full rotation from 0 to  $2\pi$ . The value of this integral is called the *total activity*  $t(\phi)$  and can be computed as  $t(\phi) = \sum_{\ell \in L} \sum_{I \in A_\phi(\ell)} |I|$ . The problem of optimizing  $t(\phi)$  is called *total activity maximization problem* (MAXTOTAL).

A valid labeling is not yet consistent in terms of the definition of Been et al. [2, 3]: while labels clearly do not jump and the labeling is independent of the rotation history, labels may still *flicker* multiple times during a full rotation from 0 to  $2\pi$ , depending on how many active ranges they have in  $\phi$ . In the most restrictive consistency model, which avoids flickering entirely, each label is either active for the full rotation  $[0, 2\pi)$  or never at all. We denote this model as *0/1-model*. In our previous paper [8] we defined a rotation labeling as consistent if each label has only a single active range, which we denote here as the *1R-model*. This immediately generalizes to the *kR-model* that allows at most  $k$  active ranges for each label. Analogously, the unrestricted model, i. e., the model without restrictions on the number of active ranges per label, is denoted as the  $\infty R$ -model.

We may apply another restriction to our consistency models, which is based on the occlusion of anchors. Among the conflicts in set  $C(\ell, \ell')$  we further distinguish *hard conflicts*, i. e., conflicts where label  $\ell$  intersects the anchor point of label  $\ell'$ . If a labeling  $\phi$  sets  $\ell$  active during a hard conflict with  $\ell'$ , the anchor of  $\ell'$  is occluded. This may be undesirable in some situation in practice, e. g., if every point in  $P$  carries useful information in the map, even if it is unlabeled. Thus we may optionally require that  $\phi(\ell, \alpha) = 0$  during any hard conflict of a label  $\ell$  with another label  $\ell'$  at angle  $\alpha$ . Note that we can include other obstacles (e. g., important landmarks on a map) which must not be occluded by a label in the form of hard conflicts. Note that a soft conflict is always a label-label conflict, while a hard conflict is always a label-point conflict (in our definition every label-point conflict induces also a label-label conflict). We showed earlier [8] that MAXTOTAL is NP-hard in the 1R-model avoiding hard conflicts and presented approximation algorithms.

### 3 Algorithmic Approaches

In this section we describe four algorithmic approaches for computing consistent active ranges that we evaluate in our experiments. We also evaluate our previous  $\frac{1}{4}$ -approximation algorithm [8] for MAXTOTAL, but omit its description due to space constraints; the full version [9] contains a sketch. Section 3.1 describes three simple greedy heuristics and Section 3.2 formulates an exact ILP model that we use primarily for evaluating the quality of the other solutions.

#### 3.1 Greedy Heuristics

In this section we describe three new greedy algorithms to construct valid and consistent labelings with high total activity. These algorithms are conceptually simple and easy to

implement, but in general we cannot give quality guarantees for the solutions computed by these algorithms.

All three greedy algorithms follow the same principle of iteratively assigning active ranges to all labels. The algorithm first initializes a set  $L'$  with all labels in  $L$ . Then it computes for each label  $\ell$  its *maximum active range*  $I_{\max}(\ell)$ , which is the active range of maximum length  $|I_{\max}(\ell)|$  such that (i)  $\ell$  is not active while in conflict with another active label that was already considered by the algorithm, and (optionally) such that (ii)  $\ell$  is not active while it has a hard conflict with another label. Initially the maximum active range of each label is either the full interval  $[0, 2\pi]$  or the largest range that avoids hard conflicts. Then the algorithm repeats the following steps. It selects and removes a label  $\ell$  from  $L'$ , assigns it the active range  $I_{\max}(\ell)$ , and updates those labels in  $L'$  whose maximum active range is affected by the assignment of  $\ell$ 's active range. If we consider the  $kR$ -model with  $k > 1$ , we keep a counter for the number of selected active ranges and add another copy of  $\ell$  with the next largest active range to  $L'$  if the counter value is less than  $k$ . The three algorithms differ only in the criterion that determines which label is selected from  $L'$  in each iteration.

The first algorithm we propose is called **GreedyMax**. In each step the algorithm selects the label with the largest maximum active range among all labels in  $L'$ . Ties are broken arbitrarily. The second algorithm, **GreedyLowCost**, determines for the maximum active range of each label the cost of adding it to the solution. This means that for each label  $\ell \in L'$  with maximum active range  $I_{\max}(\ell)$  the algorithm determines for all labels  $\ell' \in L'$  that are in conflict with  $\ell$  during  $I_{\max}(\ell)$  by how much their maximum active range would shrink. The sum of this is the *cost*  $c(\ell)$  of assigning the active range  $I_{\max}(\ell)$  to  $\ell$ . Among all labels in  $L'$  **GreedyLowCost** chooses the one with lowest cost. Finally, the last algorithm, **GreedyBestRatio** is a combination of the two preceding ones. In each step the algorithm chooses the label  $\ell$  whose ratio  $|I_{\max}(\ell)|/c(\ell)$  is maximum among all labels in  $L'$ . We conclude with a brief performance analysis of our algorithms.

**Theorem 1.** *In the  $kR$ -model with constant  $k$  the algorithm **GreedyMax** can be implemented to run in time  $O(cn \cdot (c + \log n))$  and the algorithms **GreedyLowCost** and **GreedyBestRatio** can be implemented to run in time  $O(cn \cdot (c^2 + \log n))$ , where  $n$  is the number of labels and  $c$  is the maximum number of conflicts per label in the input instance. The space consumption of all algorithms is in  $O(cn)$ .*

The running time of **GreedyMax** can be further improved to  $O(cn \log n)$ . Moreover, if all labels are unit squares, **GreedyMax** is a  $1/8$ -approximation algorithm.

### 3.2 Integer Linear Program

In this section we present an ILP-based approach to find optimal solutions for **MAXTOTAL**. This is justified since **MAXTOTAL** is NP-hard and we cannot hope for an efficient algorithm unless  $P = NP$ . We note that the same ILP formulation can also be used in the  $1/4$ -approximation algorithm to compute optimal solutions in the subinstances it considers.

The key idea of the ILP presented here is to determine regular active ranges induced by the ordered set of all conflict events. Our model contains for each label  $\ell$  and each

interval  $I$  a binary decision variable, which indicates whether or not  $\ell$  is active during  $I$ . We add constraints to ensure that (i) no two conflicting labels are active at the same time within their conflict range and (ii) at most  $k$  disjoint contiguous active ranges can be selected for each label as required in the  $k$ R-model.

*Model.* For simplicity we assume in this section that the length of each conflict range is strictly larger than 0. This assumption is not essential for our ILP formulation, but makes the description easier.

Let  $E$  be the ordered set of conflict events that also contains 0 and  $2\pi$ , and let  $E[j]$  be the interval between the  $j$ -th and the  $(j+1)$ -th element in  $E$ . We call such an interval  $E[j]$  an *atomic interval* and always consider its index  $j$  modulo  $|E| - 1$ . For each label  $\ell_i \in L$  and for each atomic interval  $E[j]$  we introduce two binary variables  $x_i^j$  and  $b_i^j$  to our model. We refer to the variables of the form  $x_i^j$  as *activity variables*. The intended meaning of  $x_i^j$  is that its value is 1 if and only if the label  $\ell_i$  is active during the  $j$ -th atomic interval; otherwise  $x_i^j$  has value 0. We use the binary variables  $b_i^j$  to indicate the start of a new active range and to restrict their total number to  $k$ . This is achieved by adding the following constraints to our model.

$$x_i^j - b_i^j \leq x_i^{j-1} \quad \forall \ell_i \in L \quad \forall j \in \{0, \dots, |E| - 2\} \quad (1)$$

$$\sum_{0 \leq j \leq |E| - 2} b_i^j \leq k \quad \forall \ell_i \in L \quad (2)$$

The effect of constraint (1) is that it is only possible to start a new active range for label  $\ell_i$  with atomic interval  $E[j]$  (i.e.,  $x_i^{j-1} = 0$  and  $x_i^j = 1$ ) if we account for that range by setting  $b_i^j = 1$ . Due to constraint (2) this can happen at most  $k$  times per label. We can also allow arbitrarily many active ranges per label as in the  $\infty$ R-model by completely omitting the variables  $b_i^j$  and the above constraints.

It remains to guarantee that no two labels can be active when they are in conflict. This can be done straightforwardly since we can compute for which atomic intervals two labels are in conflict and we ensure that not both activity variables can be set to 1. More specifically, for every pair of labels  $\ell_i, \ell_k$  and for every atomic interval  $j$  during which they are in conflict, we add the constraint

$$x_i^j + x_k^j \leq 1. \quad (3)$$

Optionally, incorporating hard conflicts can also be done easily as a hard conflict simply excludes certain atomic intervals from being part of an active range. We determine for each label all such atomic intervals in a preprocessing step and set the corresponding activity variables to 0.

Among all feasible solutions that satisfy the above constraints, we maximize the objective function  $\sum_{\ell_i \in L} \sum_{0 \leq j \leq |E| - 2} x_i^j \cdot |E[j]|$ , which is equivalent to the total activity  $t(\phi)$  of the induced labeling  $\phi$ .

This ILP considers only regular active ranges, since label activities change states only at conflict events. However, there always exists an optimal solution that is regular [8, Lemma 4], and hence we are guaranteed to find a globally optimal solution. Let  $e$  be the number of conflict events and let  $c$  be the maximum number of conflict events

per label in a MAXTOTAL instance, respectively. In the worst case the number of constraints that ensure that the solution is conflict-free (i. e., constraint (3)) is  $O(c \cdot e)$  per label, whereas we require only  $O(e)$  constraints of the other two types of constraints per label.

**Theorem 2.** *The ILP (1)–(3) solves MAXTOTAL and has at most  $O(e \cdot n)$  variables and  $O(c \cdot e \cdot n)$  constraints, where  $n$  is the number of labels,  $e$  the number of conflict events, and  $c$  the maximum number of conflicts per label.*

## 4 Experimental Evaluation

In this section we present the experimental evaluation of different labeling strategies based on the consistency models and algorithms introduced in Sections 2 and 3. We implemented our algorithms in C++ and compiled with GCC 4.7.1 using optimization level `-O3`. As ILP solver we used Gurobi 5.6. The running time experiments were performed on a single core of an AMD Opteron 2218 processor running Linux 2.6.34.10. The machine is clocked at 2.6 GHz, has 16 GiB of RAM and  $2 \times 1$  MiB of L2 cache. All reported running times are wall-clock times.

### 4.1 Benchmark Instances

Since our labeling problem is immediately motivated by dynamic mapping applications, we focus on gathering real-world data for the evaluation. As data source we used the publicly available data provided by the OpenStreetMap project ([www.osm.org](http://www.osm.org)). We extracted the latitudes, longitudes and names of *all* cities with a population of at least 50 000 for six countries (France, Germany, Italy, Japan, United Kingdom, and the United States of America) and created maps at three different scales.

To obtain a valid labeling instance several additional steps are necessary. First, the width and height of each label need to be chosen. Second, we need to map latitude and longitude to the two-dimensional plane. Third, recall that the input is a statically labeled map, and hence we need to compute such a static input labeling. For the first issue we used the same font that is used in Google Maps, i. e., `Roboto Thin`. The dimensions of each label were obtained by rendering the label’s corresponding city name in `Roboto Thin` with font size 13, computing its bounding box, and adding a small additional buffer. For obtaining two-dimensional coordinates from the latitude and longitude of each point, we used a Mercator projection (where we approximate the ellipsoid with a sphere of radius  $r = 6371\text{km}$ ). For the map scales we again wanted to be close to Google Maps. Hence, we derived instances in three different scales (65 pixel  $\hat{=}$  20km, 50km, 100km) for each country. For simplicity we refer to the scale of 65 pixel  $\hat{=}$  20km only by 20km (and likewise for the remaining scales). The last remaining step was to compute a valid input labeling. For this we used the 4P fixed-position model [6] and solved a simple ILP model to obtain a weighted maximum independent set in the label conflict graph, in which any two conflicting label positions are linked by an edge and weights are proportional to the population. Table 1 shows the characteristics of our benchmark data, which can be downloaded from [illwww.itl.kit.edu/projects/dynamiclabeling/](http://illwww.itl.kit.edu/projects/dynamiclabeling/).

Table 1: Number of labels in each benchmark instance, the number of labels in the largest connected component (lcc) and the number of connected components (cc) in the conflict graph.

<i>scales</i>	<i>countries</i>					
	FR	DE	GB	IT	JP	US
	#labels (#labels in lcc / #cc)					
20km	86 (12/51)	52 (20/26)	99 (73/19)	131 (28/48)	99 (12/34)	403 (26/203)
50km	80 (39/9)	43 (39/4)	68 (66/2)	111 (87/5)	80 (69/7)	359 (88/89)
100km	69 (69/1)	33 (33/1)	37 (37/1)	68 (68/1)	49 (44/3)	288 (213/16)

## 4.2 Evaluation of the Consistency Models

Here, we evaluate the different consistency models introduced in Section 2. The models differ by the admissible number of active ranges per label and the handling of hard conflicts. We begin by analyzing the effect of limiting the number of active ranges and consider the five models 0/1, 1R, 2R, 3R, and  $\infty$ R, all taking hard conflicts into account. As discussed in Section 2, the 0/1-model is flicker-free but expected to have a low total activity, especially in dense instances. On the other hand, the  $\infty$ R-model achieves the maximum possible total activity in any valid labeling, but is likely to produce a large number of flickering effects. Still, it serves as an upper bound on the total activities of the other models. The two most important quality criteria in our evaluation are (i) the total activity of the solution, and (ii) the average length of the active ranges.

In Table 2 we report the total activity of the optimal solution for the tested models relative to the solution in the  $\infty$ R-model. The results of the instances are aggregated by scale. We observe that the total activity of the optimal solutions in the 0/1-model drops significantly, namely to less than 55% compared to the optimal solution in the  $\infty$ R-model even for the least dense instance at scale 20km and to only 6% for a scale of 100km. Hence this model is of very little interest in practice.

We see a strong increase in the average total activity values already for the 1R-model compared to the optimal solution in the 0/1-model. For the large-scale instance 20km 1R reaches almost 95% of the  $\infty$ R-model, which has more than 19 times the number of flickering effects and active ranges of average length shorter by a factor of 1/9. For map scales of 50km and 100km, the total activities drop to 88% and 81%, respectively, but at the same time the number of flickering effects and the average active range lengths in the  $\infty$ R model are extremely poor. Thus the 1R-model achieves generally a very good labeling quality by using only one active range per label.

Finally, we take a look at the middle ground between the 1R- and the  $\infty$ R-models. It turns out that total activity of the 2R-model is off from the  $\infty$ R-model by less than 1% at scale 20km and less than 5% at scale 100km, but this increase in activity over the 1R model comes at the cost of producing twice as many flickering effects and decreasing the average active range length by 30–40%. If we allow three active ranges per label, the total activity increases to more than 99% of the upper bound in the  $\infty$ R-model at all three scales, while having significantly fewer flickering effects and longer average active ranges. The activity gain by considering the  $k$ R-model for  $k > 3$  is negligible and the disadvantage of increasing the number of flickering effects dominates.



Table 2: Average total activity of the optimal solutions with respect to the maximum possible objective value. Instances grouped by scale. Additionally we report the average interval length normalized to one full rotation, and for the  $\infty$ R-model the average number of intervals per label.

<i>model</i>	0/1		1R		2R		3R		$\infty$ R	
<i>scale</i>	total act.	total act. $\sim$ len	total act.	total act. $\sim$ len	total act.	total act. $\sim$ len	total act.	total act. $\sim$ len	$\sim$ len	$\sim$ intervals
20km	54.04%	94.56%	0.76	99.36%	0.56	99.92%	0.47	0.08	19.13	
50km	22.42%	87.79%	0.58	97.69%	0.35	99.54%	0.26	0.01	79.22	
100km	6.19%	81.01%	0.44	95.83%	0.27	99.24%	0.19	< 0.01	128.4	

We conclude that the 1R-model achieves the best compromise between total activity value and low flickering, at least for maps at larger scales with lower feature density. For dense maps the 2R- or even the 3R-models yield near-optimal activity values while still keeping the flickering relatively low. Going beyond three active ranges per label only creates more flickering but does not provide noticeable additional value.

It remains to investigate the impact of hard conflicts. For this we apply the 1R-model and compare the variant where all conflicts are treated equally (soft-conflict model) with the variant where hard conflicts are disallowed (hard-conflict model). We consider for each map scale the average relative increase in activity value of the soft-conflict model over the stricter hard-conflict model. For 20km instances the increase is 8.51%, for the intermediate scale 50km it is 19.25%, and for the small-scale map 100km the increase reaches 31.9%. These results indicate that, unsurprisingly, the soft-conflict model improves the total activity at all scales, and in particular for dense configurations of point features, where labels usually have several hard conflicts with nearby features. As discussed before, this improvement comes at the cost of temporarily occluding unlabeled but possibly important points. It is an interesting open usability question to determine user preferences for the two models and the actual effect of temporary point occlusions on the readability of dynamic maps, but such a user study is out of scope of this evaluation and left as an interesting direction for future work.

### 4.3 Evaluation of the Algorithms

In this section we evaluate the quality (total activity) and running time of the  $\frac{1}{4}$ -approximation algorithm and the three greedy heuristics GreedyMax, GreedyLowCost, and GreedyBestRatio (Section 3.1), which we abbreviate as QAPX, GM, GLC, and GBR, respectively. Additionally, we include the ILP (Section 3.2) as the only exact method in the evaluation. The ILP is also applied to optimally solve the independent subinstances considered by QAPX. In our implementation we heuristically improve the running time of the ILP by partitioning the conflict graph of the labels into its connected components and solving each connected component individually; see Table 1 for the number of labels in the largest connected component and the number of connected components in the conflict graph of each instance. For the ILP we set a time limit of 1 hour and restrict the ILP solver to a single thread. The same restrictions are applied to the ILP when solving the small subinstances in algorithm QAPX. By the design of the algorithm, a solution

obtained by QAPX will consist of many labels that have no active range, although they could be assigned one (all labels that are discarded to obtain independent cells have active range set to length 0). To overcome this drawback, we propose a combination of QAPX with the greedy algorithms. More specifically, we apply one of our greedy algorithms to each of the four solutions computed by the  $\frac{1}{4}$ -approximation and determine among the four resulting solutions the best one. In the following we refer to the combination of the  $\frac{1}{4}$ -approximation with a greedy algorithm by adding a Q in front of the greedy algorithm’s name (e. g., QGLC). We report the results of the 1R-soft-conflict model, which turned out as a reasonable compromise between low flickering and high total activity in Section 4.2.

We give a general overview of the performance of all evaluated algorithms as a scatter plot (Fig. 2). In this scatter plot each disk represents the result of an algorithm (indicated by color) applied to a single country instance. The size of the disk indicates the scale of the instance (the smaller the disk, the smaller the scale). We omitted the algorithms QGM and QGLC in this plot to increase readability, because the difference in running time and quality of the solutions between QGM, QGLC, and QGBR is negligible and creates extra overplotting.




Fig. 2: Running time (log scale) and solution quality of the algorithms in the 1R soft-conflict model.

We observe that the performance of the greedy algorithms is very good with respect to running time as well as quality of the solutions. As expected, the total activity of QAPX is always better than 25%, but generally much worse than for the remaining algorithms. It never gets close to the solutions produced by the greedy algorithms while being considerably slower. However, combining QAPX with a greedy algorithm achieves better solutions than greedy algorithms and QAPX alone, while the increase in running time over QAPX is negligible. Finally, we observe that the ILP solves the tested instances in a reasonable time frame. To obtain the optimal solution, the ILP required on average 758s, with a median of only 30.65s. However, we concede that larger instances may require significantly more time to solve, and may even be infeasible.

We now turn to a more detailed analysis of the two most promising approaches (i) using the greedy algorithms, and (ii) combining QAPX with the greedy algorithms. For a detailed depiction of the performance of the algorithms with respect to the quality of the solution see the diagrams in Fig. 3. We observe that among the three greedy algorithms GBR performs best with respect to quality with an average of 93.7%, but the difference to the other greedy algorithms is small. Even the greedy algorithm GM with the lowest total activity produces solution with an average of 91.8% of the optimal solution. Each of the combinations of QAPX with subsequent execution of a greedy algorithm outperforms each of the greedy algorithms alone in terms of quality. However, since the solutions produced by the greedy algorithms are already very close to the optimal solution, we observe only a slight increase in total activity for QGM, QGLC, and

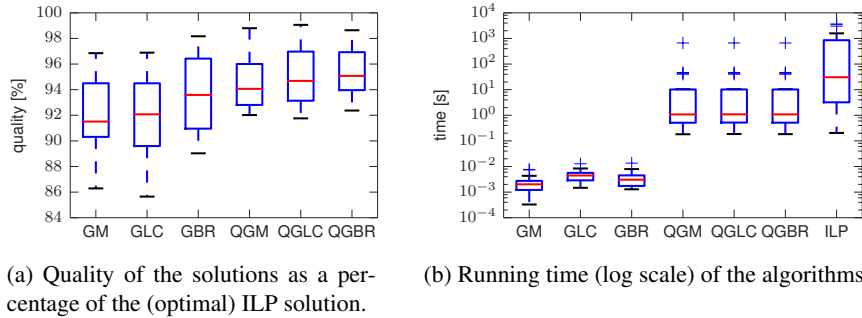


Fig. 3: Performance of the greedy algorithms and QAPX with greedy postprocessing.

QGBR over the greedy algorithms. The difference between both approaches becomes much more visible when considering the running time. While the average running time for the three greedy algorithms is between 2.5ms and 3.9ms, the average running time for the  $\frac{1}{4}$ -approximation algorithms is roughly 46s. However, we note that this large difference is mostly caused by one instance, which required over 664s to find the solution. The median running time for the enriched  $\frac{1}{4}$ -approximation algorithms is about 1.08s.

Our observations in this section are strengthened by the additional experiments (hard-conflict model and larger instances) reported in the full version of the paper [9], in which the performance of the greedy algorithms is even better. In order to give a final recommendation for an algorithm, it is necessary to make a choice on the time–quality trade-off that is acceptable in a particular situation. If running time is not the primary concern, e.g., for offline applications with high computing power available, we can recommend the ILP, which ran reasonably fast in our experiments. On the other hand, if computing power is limited and real-time labeling is necessary, e.g., on a mobile device, all three greedy heuristics can be recommended as the methods of choice; a slight advantage of GBR was observed in our experiments. All three algorithms run very fast (a few milliseconds) and empirically produce high activities of more than 90% of the optimum solution. If one wants to invest a few seconds of running time, the combination of QAPX with a greedy algorithm may be of interest as it produces slightly better solutions than the stand-alone greedy algorithms.

## 5 Conclusion

In this work, we evaluated different strategies for labeling dynamic maps that allow continuous rotation, where a labeling strategy consists of a consistency model and a labeling algorithm. In the first part of the evaluation, we considered the quality of optimal solutions in different consistency models. It turned out that the restriction to one or two active ranges per label (1R- and 2R-models) yields the best compromise in terms of low flickering and high total activity value of more than 95% of the upper bound obtained from the unrestricted model ( $\infty$ R). Additionally, treating all pairwise label conflicts as soft conflicts increased the total activity values between 8% and 32% at the cost of occasional occlusion of unlabeled point features.

In the second part of the evaluation, we investigated the performance of three new greedy heuristics and our previous  $\frac{1}{4}$ -approximation algorithm [8] in terms of labeling quality and running time. It turned out that the greedy heuristics performed very well in both total activity (well above 90%) and running time (a few ms). The unmodified  $\frac{1}{4}$ -approximation performs much worse, but the combination of  $\frac{1}{4}$ -approximation and greedy heuristics yields slightly higher total activity than the greedy heuristics alone; the running time, however, can grow to several seconds. In conclusion, we believe that the 1R model in combination with any of the three greedy algorithms is, in most cases, the best labeling strategy for labeling dynamic rotating maps. Whether the soft-conflict or the hard-conflict model is more appropriate depends on requirements of the application.

**Acknowledgment.** This work was partially supported by a Google Research Award.

## References

1. P. K. Agarwal, M. van Kreveld, and S. Suri: Label placement by maximum independent set in rectangles. *Comput. Geom. Theory Appl.*, 11(3-4):209–218 (1998).
2. K. Been, E. Daiches, and C. Yap: Dynamic map labeling. *IEEE Trans. Vis. Comput. Graph.*, 12(5):773–780 (2006).
3. K. Been, M. Nöllenburg, S.-H. Poon, and A. Wolff: Optimizing active ranges for consistent dynamic map labeling. *Comput. Geom. Theory Appl.*, 43(3):312–328 (2010).
4. M. de Berg and D.H.P. Gerrits: Approximation algorithms for free-label maximization. *Comput. Geom. Theory Appl.*, 45(4):153–168 (2011).
5. M. de Berg and D.H.P. Gerrits: Labeling moving points with a trade-off between label speed and label overlap. In: H. L. Bodlaender and G. F. Italiano (eds.) *ESA 2013*. LNCS, vol. 8125, pp. 373–384. Springer, Heidelberg (2013).
6. M. Formann and F. Wagner: A packing problem with applications to lettering of maps. In: *Proc. 7th Ann. ACM Symp. Comput. Geom.*, pp. 281–288, ACM, New York (1991).
7. A. Gemsa, B. Niedermann, and M. Nöllenburg: Trajectory-based dynamic map labeling. In: L. Cai, S.-W. Cheng, and T.-W. Lam (eds.) *ISAAC 2013*. LNCS, vol. 8283, pp. 413–423. Springer, Heidelberg (2013).
8. A. Gemsa, M. Nöllenburg, and I. Rutter: Consistent labeling of rotating maps. In: F. Dehne, J. Iacono, and J.-R. Sack (eds.) *WADS 2011*. LNCS, vol. 6844, pp. 451–462. Springer, Heidelberg (2011).
9. A. Gemsa, M. Nöllenburg, I. Rutter: Evaluation of Labeling Strategies for Rotating Maps. *CoRR*, abs/1404.1849 (2014).
10. D. S. Hochbaum and W. Maass: Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM*, 32(1):130–136 (1985).
11. M. Nöllenburg, V. Polishchuk, and M. Sysikaski: Dynamic one-sided boundary labeling. In: *Proc. 18th ACM SIGSPATIAL GIS*, pp. 310–319, ACM, New York (2010).
12. K. Ooms, W. Kellens, and V. Fack: Dynamic map labeling for users. In: *Proc. 24th Internat. Cartographic Conf.*, pp. 1–12, Military Geographic Institute, Santiago (2009).
13. M. van Kreveld, T. Strijk, and A. Wolff: Point labeling with sliding labels. *Comput. Geom. Theory Appl.*, 13(1):21–47 (1999).
14. M. Vaaraniemi, M. Treib, and R. Westermann: Temporally coherent real-time labeling of dynamic scenes. In: *Proc. 3rd Internat. Conf. Computing for Geospatial Research and Applications*, pp. 17:1–17:10, ACM, New York (2012).
15. Y. Yokosuka and K. Imai: Polynomial time algorithms for label size maximization on rotating maps. In: *Proc. 25th Canadian Conf. Comput. Geom.*, pp. 187–192, University of Waterloo, Waterloo (2013).